arXiv:1411.6011v1 [physics.flu-dyn] 21 Nov 2014

Least-biased correction of extended dynamical systems using observational data

Keith W. Myerscough[a,1,*], Jason Frank[c], Benedict Leimkuhler[b,2]

[a] *Centrum Wiskunde & Informatica, P.O. Box 94079, 1090 GB Amsterdam, the Netherlands*

[b] *School of Mathematics and Maxwell Institute for Mathematical Sciences, University of Edinburgh, James Clerk Maxwell Building, Kings Buildings Edinburgh EH9 3JZ*

[c] *Mathematical Institute, Utrecht University, P.O. Box 80010, 3508 TA Utrecht, the Netherlands*

Abstract

We consider dynamical systems evolving near an equilibrium statistical state where the interest is in modelling long term behavior that is consistent with thermodynamic constraints. We adjust the distribution using an entropy-optimizing formulation that can be computed on-the-fly, making possible partial corrections using incomplete information, for example measured data or data computed from a different model (or the same model at a different scale). We employ a thermostatting technique to sample the target distribution with the aim of capturing relavant statistical features while introducing mild dynamical perturbation (thermostats). The method is tested for a point vortex fluid model on the sphere, and we demonstrate both convergence of equilibrium quantities and the ability of the formulation to balance stationary and transient-regime errors.

Keywords: dynamical sampling, least-biased estimation, thermostat, statistical fluid dynamics, point vortex method,
2000 MSC: 65C20, 82C31, 76M23, 62B10

1. Introduction

In many applications of modern computational science the physical laws (and equations of motion) are well established yet the detailed behavior is unpredictable on long time scales due to the presence of deterministic chaos. Examples of this arise in molecular dynamics modelling [1, 2] and in the study of turbulent fluids in the atmosphere and ocean [3, 4]. For these problems, long simulations are routinely run, despite the lack of predictability, in the hope that the resulting simulation will yield useful statistical knowledge (e.g. the statistics of rare transitions between basins in molecular dynamics, or slow relaxation processes in fluids). We refer to this approach

*Corresponding author
Email addresses: k.w.myerscough@cwi.nl (Keith W. Myerscough), j.e.frank@uu.nl (Jason Frank), b.leimkuhler@ed.ac.uk (Benedict Leimkuhler)

[1] This author was supported by a grant as part of research programme 613.001.009, which is financed by the Netherlands Organisation for Scientific Research (NWO).
[2] This author was supported by grant EPSRC EP/G036136/1.

as *dynamical sampling*, where the name is suggestive of the typical requirement that simulated paths are sufficiently accurate to allow the computation of measures of dynamical mixing such as two-point temporal correlation functions [5].

In Hamiltonian systems such as molecular dynamics, it is common to run canonically prepared ensembles of microcanonical (i.e. constant energy) simulations in order to minimize the perturbation of dynamical properties. For such systems, backward error analysis [6, 7] suggests that the global behavior can best be understood not as the approximation of particular trajectory but rather as an accurate path for a perturbed continuum process described by modified equations. In the case of dynamic sampling of complex systems, the statistics of simulation data are therefore biased in that they sample an invariant measure of the modified equations, i.e. bias arises as an artifact of time discretization. Statistical bias may also arise due to spatial discretization. For example, in the setting of geophysical fluid dynamics, a comparison of discretizations of the quasi-geostrophic equations reveals that the long time mean potential vorticity field and pointwise fluctuation statistics are heavily dependent on discrete conservation laws such as energy, enstrophy, and material conservation of vorticity [8, 9, 10]. It is usually impossible to construct numerical discretizations that automatically preserve all conservation laws of statistical relevance for a given problem, so the discretization necessarily perturbs the statistical distribution. The discretization bias may be reduced by refining the discretization or by incorporating a Metropolis condition [11], but such techniques also typically lead to an increase in computational overhead, which may be unacceptable in large scale applications.

The combination of the need for computations to address both the stationary constraint ("nearness to the steady-state distribution") and to provide accuracy with respect to dynamical processes poses difficult challenges for the simulator. In this paper we consider an approach to perturbing dynamics to correct statistical bias in systems at statistical equilibrium. If the statistical distribution is completely specified via a probability density function (pdf) it can be sampled using a "thermostat." Such thermostats, originating in molecular dynamics, can be extended to handle both smooth [12] and nonsmooth [13] densities and to treat noncanonical Hamiltonian systems. In [14], a thermostat was used as a model reduction technique for a vortex model of a fluid (suppressing the detailed interactions of a few strong vortices with a weak vortex field). In another recent article [15], thermostats have been suggested as a means of sampling incompletely specified systems (with noisy gradients), with applications in learning theory. The standard framework of thermostating used in these and other applications assumes a fixed, known distribution such as the Gibbs-Boltzmann distribution. In this article, we assume that, instead of the pdf, what is available is a partial set of *expectations* of observables with respect to the unknown invariant measure, which may arise from experiment or other types of modelling. In this setting, information theory (in particular entropy maximization [16, 17]) offers tools for constructing least-biased densities, close to some known prior distribution, which are consistent with observations. The iterative method (based on [18, 19, 20]) involves computation of Lagrange multipliers (one for each observable) that modify the probability density. The Lagrange multipliers are computed using an iterative procedure in which each stage represents an ensemble average (with respect to the previous estimate of the density). To make the method practical in situations where the sampling is costly, we consider an adaptive procedure which uses only short-time ensemble bursts to gradually tune the parameters in simulation. At the same time, we are able to show in numerical experiments that autocorrelation functions are only modestly perturbed meaning that we would expect to be able to recover dynamical information such as diffusion and other transport coefficients.

We emphasize that the framework of least-biased estimation is well known but applied here in

a novel way. A related technique is used by Majda and Gershgorin [21] to develop a framework for validating computational models and choosing the optimal linear combination of an ensemble of model outputs, so as to minimize the discrepancy between the ensemble distribution and the least-biased estimate, arguing that the latter is the best available measure for comparison, when the true invariant measure is unknown. With the approach we develop here, we enforce exact adherence to the least-biased measure, which is constructed automatically in simulation, by perturbing the dynamics to take full advantage of available information.

The remainder of this article is organized as follows. In the next section, we discuss the maximum entropy framework for correcting the density to reflect thermodynamic constraints. We apply and evaluate the method in the setting of a system of point vortices on the surface of a sphere, which represents a simple geophysical model with multiple statistically relevant first integrals.

2. Bias correction method

Our interest is in extended dynamical systems with many degrees of freedom that evolve near statistical equilibrium. Further, we imagine that we are given a simplified dynamical model for the evolution of some projection (i.e. a "coarse graining") of the phase variables (coarse grained variables $y(t) \in \mathbf{R}^d$) . Although the original system is complex and its details unknown, we assume that we can obtain in some way (e.g. through measurement) a collection of "observations" of mean values of functions of the reduced variables. That is there are functions $C_k : \mathbf{R}^d \to \mathbf{R}$, $k = 1, 2 \ldots, K$ and given values c_k, $k = 1, 2, \ldots, K$, such that

$$c_k = \langle C_k(y) \rangle, \quad k = 1, \ldots, K, \tag{1}$$

where $\langle C_k(y) \rangle$ represents averaging with respect to the true, empirical invariant measure of the dynamical system. Our goal is to find a perturbed dynamical model for the reduced variables which (a) is compatible with the indicated thermodynamic constraints (1), and (b) weakly perturbs the dynamics compared to those of the native model.

Empirical information theory generalizes the principle of insufficient reason, by proposing the least-biased probability density consistent with a set of observations. See the classical work of Jaynes [16, 17], the monographs [19, 22] and an extensive treatment in the geophysical fluid context in the monograph by Majda and Wang [23]. The least-biased density is defined as the probability density $\rho(y)$ that maximizes the information entropy functional

$$\mathcal{S}[\rho] = -\int_{\mathcal{D}} \rho(y) \log \rho(y)\, dy,$$

subject to a set of constraints given by observations. When $\mathcal{D}$ is a compact set and there are no observations, the minimizer is the uniform density $\rho \equiv |\mathcal{D}|^{-1}$. The entropy $\mathcal{S}$ is the unique measure of uncertainty that is positive valued, monotonically increasing as a function of uncertainty, and additive for independent random variables. With observable functions $\{C_k(y) | k = 1, 2, \ldots K\}$ let

$$\mathbb{E}_\rho C_k = \int_{\mathcal{D}} C_k(y) \rho(y)\, dy \tag{2}$$

denote expectation in the (as yet undetermined) density ρ. Defining Lagrange multipliers λ_k, $k = 1, \ldots, K$, associated with the observables C_k, the constrained minimization problem is

$$\rho = \arg\max_{\hat{\rho}} \left[\mathcal{S}[\hat{\rho}] - \sum_{k=1}^{K} \lambda_k \left(\mathbb{E}_{\hat{\rho}} C_k(y) - c_k \right) \right].$$

When it exists, the maximum entropy solution satisfies

$$\rho(y) = \lambda_0 \exp\left(-\lambda_1 C_1(y) - \cdots - \lambda_K C_K(y)\right),$$

where λ_0 is chosen to satisfy $\int_{\mathcal{D}} \rho \, dy = 1$, and λ_k is chosen such that $\mathbb{E}_\rho C_k(y) = c_k$.

In some cases, besides the observations, we may be given prior statistical information on the process $y(t)$. The Kullback-Leibler divergence, or relative entropy,

$$\mathcal{S}[\rho(y)] = \int \rho(y) \ln \frac{\rho(y)}{\pi(y)} \, dy$$

which represents a (non-symmetric) distance between measures. It quantifies the information lost in approximating $\rho(y)$ by $\pi(y)$.

Suppose y is a random variable with distribution (law) $y \sim \rho$, where ρ is unknown. Suppose further, that we are given a prior distribution π, presumed to be close to ρ, and a set of K observations (1). Following Jaynes [16, 17], the least-biased distribution ρ consistent with the observations c_k and prior π solves the constrained minimization problem

$$\rho = \arg\min_\rho \left[\mathcal{S} - \lambda_0 \left(1 - \int \rho(y) \, dy \right) - \sum_{k=0}^{K} \lambda_k \left(c_k - \int C_k(y) \rho(y) \, dy \right) \right],$$

where the λ_k are Lagrange multipliers to enforce the condition that the expectations (2) agree with the observations (1). The solution to the variational problem is

$$\rho(y) = \lambda_0 \exp\left(-\lambda_1 C_1(y) - \cdots - \lambda_K C_K(y)\right) \pi(y), \tag{3}$$

where the Lagrange multipliers λ_k are chosen consistently with the observations (1) and λ_0 is a normalization constant so that ρ is a probability density function.

Methods for determining the Lagrange multipliers are discussed in [18, 19, 20]. We use the following algorithm based on re-weighting. Assume we are given a sequence of samples y^n, $n = 1, \ldots, N$, distributed according to a known prior distribution $\pi(y)$, i.e. $y^n \sim \pi$. The expectation under $\pi(y)$ of a function $\Phi(y)$ has the consistent and unbiased estimator

$$\widehat{\Phi}^\pi = \frac{1}{N} \sum_{n=1}^{N} \Phi(y^n).$$

Given the posterior distribution $\rho(y)$ of the form (3), compute the expectation $\mathbb{E}_\rho \Phi$ by re-weighting of the integral

$$\mathbb{E}_\rho \Phi = \int \Phi(y) \rho(y) \, dy = \lambda_0 \int \Phi(y) e^{-\sum_{i=1}^K \lambda_i C_i(y)} \pi(y) \, dy = \lambda_0 \mathbb{E}_\pi \{ \Phi(y) \lambda_0 e^{-\sum_{i=1}^K \lambda_i C_i(y)} \},$$

yielding an unbiased estimator for $\mathbb{E}_\rho \Phi$ given by

$$\widehat{\Phi}^\rho = \frac{\lambda_0}{N} \sum_{n=1}^{N} \Phi(y^n) e^{-\sum_{i=1}^K \lambda_i C_i(y^n)}.$$

We wish to ensure that the observations c_k satisfy

$$c_k = \widehat{C}_k^\rho = \frac{\lambda_0}{N} \sum_{n=1}^{N} C_k(y^n) e^{-\sum_{i=1}^K \lambda_i C_i(y^n)}, \quad k = 1, \ldots, K.$$

We can use this fact to define a Newton-Raphson iteration to determine the Lagrange multipliers λ_k. Define the residual r with components

$$r_k(\lambda) = c_k - \frac{\lambda_0}{N}\sum_{n=1}^{N} C_k(y^n) e^{-\sum_{i=1}^{K}\lambda_i C_i(y^n)}, \quad k = 1, \ldots, K,$$

with $\lambda = (\lambda_1, \ldots, \lambda_K)$ and $r = (r_1(\lambda), \ldots, r_K(\lambda))$. Note that λ_0 can be viewed as a function of $\lambda_1, \lambda_2, \ldots, \lambda_K$ chosen from the normalization condition, i.e.,

$$\lambda_0 = \left[\sum_{n=1}^{N} e^{-\sum_{j=1}^{K}\lambda_j C_j(y^n)}\right]^{-1}.$$

The Jacobian matrix $J = (J_{kj})$ of the vector function r is determined as

$$J_{kj}(\lambda) := \frac{\partial r_k}{\partial \lambda_j} = \frac{\lambda_0}{N}\sum_{n=1}^{N} C_k(y^n) C_j(y^n) e^{-\sum_{i=1}^{K}\lambda_i C_i(y^n)} \quad j, k = 1, \ldots, K.$$

The iteration then proceeds as $\lambda^{\alpha+1} \leftarrow \lambda^\alpha - J^{-1}(\lambda^\alpha) r(\lambda^\alpha)$.

2.1. Adaptive determination of Lagrange multipliers

In many cases it will be difficult or costly to carry out a complete sampling of the distribution at each iteration step of the Newton procedure. Moreover, the standard framework excludes applications where (i) the statistical knowledge is expected to improve as the simulation progresses, (ii) the average observables are known to vary slowly with time, or (iii) it is unfeasible to constuct a large enough ensemble distributed in the prior. For these cases we consider using the simulation data of a small ensemble (propagated in short bursts of M timesteps) for updating the Lagrange multipliers for mean observation data. This results in an adaptive algorithm for obtaining the Lagrange multipliers "on-the-fly" during simulation.

Consider the following: an ensemble of P simulations (preferably with initial conditions distributed close to $\pi(y)$) is advected $M\Delta t$ in time, where M is chosen sufficiently large such that the ensemble members sample π well. These ensemble members can be used in an estimator for $\mathbb{E}_{\lambda^1} C_k$ given by

$$\widehat{C}_k^{(1)} = \frac{1}{P}\sum_{p=1}^{P} C_k(y_0^p)\lambda_0^1 \exp\left(-\sum_{i=1}^{K}\lambda_i^1 C_i(y_0^p)\right), \quad k = 1, \ldots, K,$$

where the superscript (1) indicates that it is an estimator for a distribution with Lagrange multipliers λ_j^1. A Newton-Raphson iteration to find the first set of Lagrange multipliers such that observations match data has the residual

$$r_k^1 = \widehat{C}_k^{(1)} - c_k = \frac{1}{P}\sum_{p=1}^{P} C_k(y_0^p)\lambda_0^1 \exp\left(-\sum_{i=1}^{K}\lambda_i^1 C_i(y_0^p)\right) - c_k, \quad k = 1, \ldots, K.$$

Using this updated value for λ the simulations will sample the distribution $\rho \propto e^{-\sum_{i=1}^{K}\lambda_i^1 c_i(y)}$ after some time $M\Delta t$. (See below for some practical issues associated to this.) Using these

samples alongside the initial data, λ_j^2 is found. Iteration of this process leads to the following equation

$$r_k^m = \widehat{C}_k^{(m)} - c_k = \frac{1}{mP}\sum_{p=1}^{P}\sum_{\ell=0}^{m-1} C_k(y_{\ell M}^p)\lambda_0^\ell \exp\left(\sum_{i=1}^{K}(\lambda_i^l - \lambda_i^m)C_j(y_{\ell M}^p)\right) - c_k,\ k = 1,\ldots,K, \quad (4)$$

where we remind that, at each stage of iteration, λ_0^l is a function of the multiplier vector λ^l (indices $1\ldots K$). In the calculation (4) $\lambda_1^0,\ldots\lambda_K^0$ would ideally be zero. There are cases where it is impossible to obtain an accurate initial distribution according to the prior, in which case the initial Lagrange multipliers can be chosen different from zero if initial conditions sampling $\pi\lambda_0^0\exp\left(-\sum_k \lambda_k^0 C_k\right)$ are easier to find than those sampling just the prior π. Solutions of (4) are found using Newton-Raphson iteration. The gradient is given by

$$\frac{\partial r_k^m}{\partial \lambda_j^m} = \frac{1}{mP}\sum_{p=1}^{P}\sum_{\ell=0}^{m-1} C_k(y_{\ell M}^p)C_j(y_{\ell M}^p)\lambda_0^\ell \exp\left(\sum_{i=1}^{K}(\lambda_i^l - \lambda_i^m)C_i(y_{\ell M}^p)\right), \quad j,k = 1,\ldots,K. \quad (5)$$

In this way the Lagrange multipliers may be found "on-the-fly."

As a convergence result let us consider the case where both P and M may be chose arbitrarily large. For $P \to \infty$ the Lagrange multipliers computing using only the initial data sampling the prior will be correct. Given sufficiently large M the samples after evolving the thermsotated system $M\delta t$ in time will accurately sample the distribution corresponding to these Lagrange multipliers. The ensemble averages will then correspond to the observations, and the Lagrange multipliers no longer need updating.

2.1.1. Adaptive algorithm

There are two important practical modifications to the algorithm described above that are included in the numerical implementation of this method:

- The first modification is limiting the rate of change of the Lagrange multipliers. If the Lagrange multipliers change rapidly, the thermostat may require a long time to equilibriate. This requires a larger value for M, increasing the simulation time required before including new samples. The effect is especially noticable at the beginning of a simulation, due to two factors: (i) the small sample size leads to inaccurate expectations for the observables, and (ii) the Lagrange multipliers may be far from their correct value. By limiting the rate of change of the Lagrange multipliers, these problems are circumvented.
- The second modification regards the number of samples included when updating the Lagrange multipliers. In equations (4) and (5) *all* previous values λ_k^ℓ are included. In a long simulation, this leads to a growing computational demand. By taking only a fixed number (q) of recent steps the computational demand can be reduced. In the case that the initial samples cannot accurately be drawn from the prior, this has the further advantage that these inaccuracies are eventually forgotten.

The algorithm, including these practical modifications, is summarized in Algorithm 1.

Given initial conditions according to prior $\pi(y)$
Set initial Lagrange multipliers to zero
for $m \leftarrow 1$ **to** n **do**
  for $j \leftarrow 1$ **to** M **do**
    advance simulation one time step using current value for the Lagrange multipliers
  end
  store relevant simulation observables for time step im.
  while $|\langle C_k(y)\rangle_{\lambda} - c_k| >$ tolerance **do**
    compute residual using reweighted samples at times $M \times \max(m-q,0), \ldots, mM$
    compute residual gradient using the same data
    update Lagrange multiplier estimation
  end
  limit the change in Lagrange muliplier (if necessary)
end

Algorithm 1: Adaptive determination of Lagrange multipliers "on-the-fly"

2.2. Thermostat

We introduce the bias-correction methodology for a Hamiltonian dynamical system

$$\frac{dy}{dt} = f(y) = B(y)\nabla H(y), \qquad y(t) \in \mathcal{D},\ B(y) = -B(y)^T,\ H(y) : \mathcal{D} \to \mathbf{R}, \tag{6}$$

possessing a divergence-free vector field[3] $\nabla \cdot f \equiv 0$. Invariance of the Hamiltonian H along solutions of (6) follows from $\frac{d}{dt}H(y(t))) = \nabla H \cdot \frac{dy}{dt} = \nabla H \cdot B \nabla H = 0$, due to skew-symmetry of $B(y)$. Additional first integrals are often present: $I_\ell(y) : \nabla I_\ell \cdot f \equiv 0$, $\ell = 1, \ldots, L$. In this paper we consider only the case where all observables of the physical process $\mathcal{Y}(t)$ correspond to functions of the conserved quantities $\{H, I_\ell, \ell = 1, \ldots, L\}$, that is, $C_k(y) = C_k(H(y), I_1(y), \ldots, I_L(y))$, $k = 1, \ldots, K$.

Thermostats are used in molecular dynamics to model the trajectories of molecules in a fluid at constant temperature. From statistical mechanics, it is well known that the trajectories of a system of particles in thermal equilibrium with a reservoir at constant temperature sample the canonical or Gibbs distribution, which has global support. The governing equations are Hamiltonian, however, implying that the trajectories are restricted to a level set of Lebesgue measure zero. Hence, to model a system at constant temperature, it is necessary to perturb the vector field to make trajectories ergodic with respect to the Gibbs distribution. The most common way of achieving this is by adding suitable stochastic and dissipative terms satisfying a fluctuation-dissipation relation (Langevin dynamics). An advantage of Langevin dynamics is provable ergodicity with respect to the Gibbs distribution [24]. However, DelSole [25] warns that direct stochastic forcing of trajectories leads to inaccurate dynamical quantities since auto-correlation functions are strongly perturbed. For smooth deterministic Hamiltonian dynamics, normalized velocity autocorrelation functions are of the form $1 - c\tau^2$, $c > 0$ in the zero-lag limit $\tau \to 0$, whereas the autocorrelation of a variable that is directly forced by white noise must take the form $\exp(-\kappa\tau)$, $\kappa > 0$ in the same limit. This implies that the direct stochastic perturbation leads to auto-correlation functions that have nonzero slope and opposite curvature at zero lag.

[3]The latter condition is automatic for systems (6) with constant B. Strictly speaking, the approach described here is applicable to any system with divergence-free vector field $\nabla \cdot f \equiv 0$ possessing one or more first integrals.

An alternative approach, pioneered by Nosé [26, 27] and Hoover [28] proceeds to augment the phase space by one dimension through coupling of (6) to an additional thermostat variable $\xi(t)$. The dynamics of ξ are constructed to ensure that the extended dynamics on $\mathbf{R}^{d+1}$ preserves an equilibrium density whose marginal on $\mathbf{R}^d$ is the target (e.g. Gibbs) density. The fully deterministic thermostats of Nosé and Hoover have no mechanism to guarantee ergodicity with respect to the target density, and hence have been modified by various authors who include stochastic forcing of the thermostat variable ξ, leading to the so-called Nosé-Hoover-Langevin method [29, 30, 31]. A generalization to generic Hamiltonian systems is the Generalized Bulgac-Kusnezov (GBK) [32, 33] thermostat:

$$dy = f(y)dt + \xi\varepsilon g(y)dt \tag{7a}$$

$$d\xi = \varepsilon h(y)dt - \gamma\xi dt + \sqrt{2\gamma}dw, \tag{7b}$$

where $\varepsilon > 0$ and $\gamma > 0$ are parameters, $w(t)$ is a scalar Wiener process, and g and h are discussed below. Given a target density $\rho(y) \propto \exp(-A(y))$, $A : \mathcal{D} \to \mathbf{R}$, denote the augmented product density by $\tilde{\rho}(y, \xi) = \rho(y) \cdot \mu(\xi)$, with μ a univariate normal distribution with mean zero and standard deviation one. It is easily checked that $\tilde{\rho}$ is stationary under the Fokker-Planck operator associated with (7) provided

$$h(y) = \nabla \cdot g - g \cdot \nabla A. \tag{8}$$

Furthermore it is argued in [13] that the target measure is ergodic provided the vector fields f and g satisfy a Hörmander condition. In some cases it is also desirable to use the freedom in choosing g to ensure preservation of some first integrals of the vector field f. We will see an example of this later in this paper.

The parameter ε can be used to control the relative strength of the thermostat compared to that of the unperturbed vector field f. This will affect the rate at which the invariant measure is sampled, but has no influence on the measure itself. It has been proved in [31] and observed numerically in [34, 13] that GBK/NHL thermostating leads to a weak perturbation of the original trajectories in the sense that autocorrelation functions preserve the leading terms, i.e. have the form $1 - c\tau^2 + O(\tau^3)$, as $\tau \to 0$. The GBK thermostat is applicable when the vector field f is divergence free $\nabla \cdot f \equiv 0$ and when the target density ρ is a function of first integrals of f.

We therefore propose (1) constructing a least-biased information theoretic target density based on observations of functions of conserved quantities (with or without prior distribution), followed by (2) thermostated perturbation of dynamics to ensure sampling of the target distribution with a GBK thermostat. The thermostating method is incorporated into Algorithm 1 to provide the scheme for sampling the adapted, data-dependent distribution.

3. Application to reduced modelling of point vortices

In this section we apply the least-biased correction methodology to the simple model of point vortices on the sphere. We choose this model because it has a Poisson structure and multiple conserved quantities, including total energy and angular momentum and a set of Casimirs, with various degrees of statistical significance. Although one can construct a point vortex approximation of quasi-geostrophic potential vorticity dynamics, we ignore the effects of topography and finite deformation radius for computational simplicity.

3.1. Point vortex system

A simple conceptual model of the atmosphere is given by the quasigeostrophic potential vorticity equation on a rotating sphere:

$$\frac{Dq}{Dt} \equiv \frac{\partial q}{\partial t} + u \cdot \operatorname{grad} q = 0, \quad \operatorname{div} u = 0, \quad q = \operatorname{curl} u + f_0 + h, \tag{9}$$

where q denotes the potential vorticity, u is the velocity field in the tangent plane, assumed divergence-free, $f_0 = 2\Omega \sin\theta$ is the local Coriolis force, and h is the surface topography.

The vortex approximation of (9) is well known. For algorithms and analysis of the dynamics of point vortex systems, see the books [35, 36]. For advanced modelling and convergence analysis in the continuum limit, see [37]. For numerical computation with point vortices, it is advantageous to embed the sphere in $\mathbf{R}^3$. In the sequel we will denote vectors in $\mathbf{R}^3$ by bold type. For simplicity we neglect topography, taking $h \equiv 0$, under which assumption the quasigeostrophic model is equivalent to the 2D Euler equations. We may then also ignore rotation (i.e. $f_0 \equiv 0$) as it gives rise to a trivial rigid body rotation of the ensuing point vortex system [35]. Since the velocity field is divergence-free in the tangent plane, it can be represented in terms of a stream function ψ as

$$\boldsymbol{u} = \hat{\boldsymbol{k}} \times \nabla\psi$$

where $\hat{\boldsymbol{k}}$ is the unit normal vector on the surface of the sphere. The potential vorticity and stream function are related by $\Delta\psi = q - f_0 - h$ with Δ the Laplace-Beltrami operator (from which it is apparent that topography, if included, would lead to a nonhomogeneous background term in the stream function).

A point vortex system is constructed by taking the vorticity field in (9) to be a sum of Dirac distributions

$$q(x,t) = \sum_{i=1}^{M} \Gamma_i \delta(\boldsymbol{x} - \boldsymbol{x}_i),$$

where Γ_i is the vortex strength or circulation of the ith point vortex. The point vortices induce a stream function $\psi(\boldsymbol{x}) = \sum_i \frac{-1}{4\pi} \sum \Gamma_i \ln(2 - 2\boldsymbol{x} \cdot \boldsymbol{x}_i(t))$ as a sum of Green's functions of the Laplacian. The unit normal on the sphere is given by $\hat{\boldsymbol{k}} = \boldsymbol{x}/|\boldsymbol{x}|$. Because vorticity is materially conserved in the velocity field, the motion of point vortices is given by $\dot{\boldsymbol{x}}_i = \boldsymbol{u}(\boldsymbol{x}_i)$, i.e.,

$$\dot{\boldsymbol{x}}_i = \boldsymbol{x}_i \times \nabla\psi(\boldsymbol{x}_i) \quad i = 1, 2, \ldots, M,$$

where a unit sphere will assumed. The equations of motion may also be written as a Hamiltonian system with Lie-Poisson structure

$$\Gamma_i \dot{\boldsymbol{x}}_i = \boldsymbol{x}_i \times \nabla_{\boldsymbol{x}_i} H \quad i = 1, 2, \ldots, M, \tag{10}$$

where the Hamiltonian, defined by $H = \int_{\mathcal{D}} \frac{1}{2}|\boldsymbol{u}|^2 d\boldsymbol{x}$, is given by

$$H = -\sum_{i=1}^{M} \sum_{j=1}^{i-1} \frac{\Gamma_i \Gamma_j}{4\pi} \ln(2 - 2\boldsymbol{x}_i \cdot \boldsymbol{x}_j).$$

By introducing $y = \left(\boldsymbol{x}_1^T, \boldsymbol{x}_2^T, \ldots, \boldsymbol{x}_M^T\right)^T$, equation (10) can be written in the more compact form

(6) with the block-diagonal structure matrix

$$B(y) = \begin{bmatrix} \Gamma_1^{-1}\widehat{\boldsymbol{x}}_1 & & & \\ & \Gamma_2^{-1}\widehat{\boldsymbol{x}}_2 & & \\ & & \ddots & \\ & & & \Gamma_M^{-1}\widehat{\boldsymbol{x}}_M \end{bmatrix},$$

where $\widehat{\boldsymbol{x}}_i$ denotes the 3×3 skew-matrix satisfying $\widehat{\boldsymbol{x}}_i \boldsymbol{a} := \boldsymbol{x}_i \times \boldsymbol{a}$, for all $\boldsymbol{a} \in \mathbf{R}^3$. The Poisson bracket for the system is given equivalently by

$$\{F, G\} = \sum_{i=1}^{M} \frac{1}{\Gamma_i} \nabla_{\boldsymbol{x}_i} F \cdot (\boldsymbol{x}_i \times \nabla_{\boldsymbol{x}_i} G) \quad \text{or} \quad \{F, G\} = \nabla F(y)^T B(y) \nabla G(y).$$

This Poisson structure is a generalization of the rigid body Poisson structure and also occurs in ferromagnetic spin lattices [38, 39, 40] and elastic rods (e.g. [41]).

The vortex positions are defined in Cartesian coordinates, but initial positions $\boldsymbol{x}_i(0)$ are chosen on the sphere. Because each $|\boldsymbol{x}_i|$ is a Casimir of the Poisson bracket it is ensured that the vortices remain on the sphere. This restricts the effective phase space of the system to the direct product of M spheres S^2. Furthermore, the rotational symmetry of the sphere gives rise to three Noether momenta, which are expressed by the angular momentum vector

$$\boldsymbol{J} = \sum_{i=1}^{M} \Gamma_i \boldsymbol{x}_i. \tag{11}$$

When studying the statistics of point vortices on the disk, Bühler [42] did not observe the (planar) angular momentum to be of great importance. On the sphere, however, the angular momentum *does* play an important role in the statistics.

The GBK thermostat (7) is only applicable to nondivergent systems $\nabla \cdot f \equiv 0$. It is straightforward to check that this condition holds for the spherical point vortex model.

3.2. Time integration

A numerical integrator can be constructed by splitting the differential equations into integrable subproblems (see related ideas in [43, 44]). We develop a new integrator for the system in Appendix A which exactly preserves all Casimir functions of the system. Furthermore, backward error analysis for symplectic integrators can be extended to Poisson systems to explain approximate conservation of the Hamiltonian [6]. Due to the additive form of the angular momentum vector (11), it may also be preserved exactly using a pairwise splitting. By expanding the Hamiltonian into its pairwise terms in the dynamics we find

$$\dot{y} = B(y)\nabla H(y) = \sum_{i<j} B(y)\nabla H_{ij}(\boldsymbol{x}_i, \boldsymbol{x}_j), \qquad H_{ij} = \frac{\Gamma_i \Gamma_j}{8\pi} \ln\left(2 - 2\boldsymbol{x}_i \cdot \boldsymbol{x}_j\right) \tag{12}$$

Each pairwise interaction is represented by the dynamical system $\dot{y} = B\nabla H_{ij}$ with the associated time-Δt flow map $\phi_{\Delta t}^{i,j}$. The time-Δt flow map of the dynamics $B\nabla H$ may be approximated by a symmetric composition of pair flows

$$\Phi_{\Delta t} = \prod_{(i,j)\in C} \phi_{\Delta t/2}^{i,j} \circ \prod_{(i,j)\in C^*} \phi_{\Delta t/2}^{i,j}, \tag{13}$$

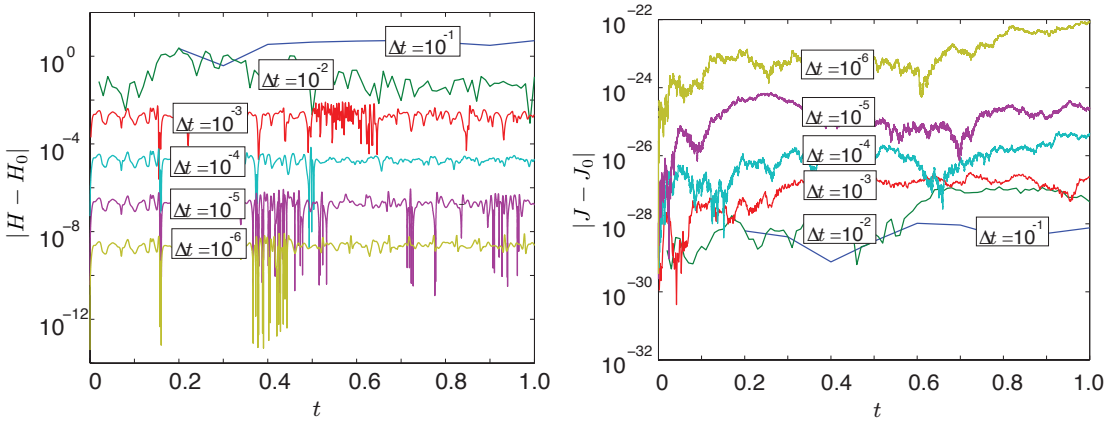

Figure 1: Error in the energy (left) and angular momentum (right) for different time step sizes. Both from simulations with 16 vortices, of which 4 with strength ± 5 and the rest with strength ± 1. The momentum error is due to limited floating point accuracy. With decreasing time step the number of time steps increases and the inaccuracies accumulate, but they remain well within any reasonable demand for accuracy.

where C is an ordered set of all possible pairs (i, j) with $i < j$ and C^* denotes the reverse ordering. This symmetric splitting yields a consistent numerical method of second order accuracy. The details of the integration procedure involving exact solution of the pairwise interaction is detailed in Appendix A.

Because the flow map of each vortex pair is the exact solution of the local Poisson system $\dot{y} = B(y)\nabla H_{ij}$ and also respects the Casimirs of the system, the composition $\Phi_{\Delta t}$ is a Poisson integrator [6, p. 247]. Expanding the angular momentum vector as $\boldsymbol{J} = \boldsymbol{J}_{ij} + \sum_{k \neq i,j} \Gamma_k \boldsymbol{x}_k$ we note that the time integration of any pair (i, j) preserves the local angular momentum $\boldsymbol{J}_{ij}$ and leaves the other vortices untouched. Hence the angular momentum is exactly conserved by the splitting method. The Hamiltonian is not exactly conserved under the motion of vortex pair, but the error can be studied by backward error analysis; see e.g. [6]. Figure 1a shows the error in the energy for simulations over a range of time step sizes, confirming second order convergence. The angular momentum should be conserved exactly by the Strang splitting. The results displayed in Figure 1b confirm this as the errors are always well within machine precision.

3.3. Thermostat perturbation vector

There is flexibility in the choice of the perturbation vector field $g(y)$ in (7). Its selection is determined both by the need for ergodicity with respect to the target measure and the need to preserve some invariants of the vector field $f(y)$. We distinguish between invariants of f whose values are known precisely, due to problem geometry for instance, and those whose values are uncertain and only known in expectation. For point vortices on the sphere, the lengths of the vortex positions $|\boldsymbol{x}_i|$ are Casimir invariants, arising from the embedding of the sphere in $\mathbf{R}^3$, and are not subject to uncertainty. We choose a perturbation vector $g(y)$ that respects these structural invariants.

The double-bracket dissipation developed in [45] preserves Casimirs of the original system and

is a candidate for $g(y)$:

$$\tilde{g}_i(\boldsymbol{x}_i) = \sum_{j \neq i} \boldsymbol{x}_i \times \boldsymbol{x}_i \times \frac{\Gamma_j}{4\pi} \frac{\boldsymbol{x}_j}{1 - \boldsymbol{x}_i \cdot \boldsymbol{x}_j}. \tag{14}$$

The denominator in (14) causes stiffness when like-signed vortices approach one another, restricting the step size of an explicit splitting method. To alleviate these matters we use a modified scheme defined by

$$g_i(\boldsymbol{x}_i) = \sum_{j \neq i} \boldsymbol{x}_i \times \boldsymbol{x}_i \times \frac{\Gamma_j}{4\pi} \boldsymbol{x}_j. \tag{15}$$

The desirable properties of the thermostat are unaffected by this modification. Appendix B contains a detailed description of the numerical integration of these dynamics.

The thermostat (7) is designed to sample a target density $\rho(y) \propto e^{-A(y)}$ on the phase space of y. The thermostat variable ξ is normally distributed, yielding the extended distribution $\rho \propto e^{-A(y)-\frac{1}{2}\xi^2}$. The perturbation vector field g must additionally ensure that the thermostated system is ergodic in the target density. Because the target measure is positive for all open sets on the phase space, hypoellipticity of the Fokker-Planck equation associated with (7) is sufficient to prove uniqueness of the invariant measure [13]. Hypoellipticity follows from Hörmander's controllability condition [46]. The condition has been tailored to GBK thermostats in [13], but it is difficult to check in practice. Here we instead check empirically that single trajectories have statistics that agree with the target distribution.

3.4. Maximum entropy model

To apply the methodology proposed in Section 2 in the setting of a reduced model for point vortices, we use point vortices distributed evenly over the surface of the sphere as the prior π. it remains to specify for which functions of the conserved quantities H and $\boldsymbol{J}$ the expectations will be observed during simulation of the full model.

In [14] a thermostat was used to model a set of point vortices on a disk in the canonical ensemble. To accurately reproduce statistics from a full model with a moderate number of point vortices, it was necessary to modify the canonical density with a term quadratic in the Hamiltonian, that is, a density of the form $\rho(y) \propto \exp(-\beta H(y) - \gamma H(y)^2)$. Motivated by the experience in [14], we choose observations that include linear and quadratic functions of H and $\boldsymbol{J}$.

If the angular momentum of the full system is zero, then there is no directional preference for the angular momentum vector $\boldsymbol{J}$. We consider the following set of observables:

$$C_1 = H, \quad C_2 = |\boldsymbol{J}|^2, \quad C_3 = H^2, \quad C_4 = |\boldsymbol{J}|^4, \quad C_5 = H^2|\boldsymbol{J}|^2, \tag{16}$$

and denote the corresponding Lagrange multipliers by β_H, β_J, γ_H, γ_J, γ_{HJ}.

The least-biased density consistent with observations of the $\mathbb{E}C_k$ is

$$\tilde{\rho}(H) = e^{-\beta_H H - \beta_J |\boldsymbol{J}|^2 - \gamma_H H^2 - \gamma_J |\boldsymbol{J}|^4 - \gamma_{HJ} H |\boldsymbol{J}|^2}. \tag{17}$$

4. Numerical comparison

To verify the methodology proposed in this article for correcting expectations, we apply it to a reduced model of point vortices similar to the configuration used in [47, 14]. We distinguish

between three models. The *full model* consists of a system (10) of $M_{\text{full}} = 288$ point vortices, of which 8 strong vortices of circulation $\Gamma_j = \pm 1$ and 280 weak vortices of circulation $\Gamma_j = \pm\frac{1}{5}$. Both strong and weak classes are comprised of equal numbers of positively and negatively oriented point vortices. The *reduced model* consists of (10) with just $M = 8$ strong vortices. Finally, the *corrected model* consists of a thermostated system (7) with unperturbed vector field f given by (10) for $M = 8$ strong vortices, perturbation vector field g given by (15), and equilibrium measure defined by the least-biased density (17). Additionally, we compare with Metropolis-Hastings samples from the least-biased density (17) to help distinguish between errors incurred due to the maximum-entropy model and those due to the thermostat.

We run seven long simulations of the full model with angular momentum vector $\boldsymbol{J}_{\text{full}} = \mathbf{0}$ and total energies chosen from the set $H_{\text{full}} \in \{-2, -1, 0, 1, 2\}$. For each run we determine the time averages of the observables (16) for the subset of strong vortices. When computing the Hamiltonian H we include only the internal coupling between strong vortices. The time averages are tabulated in Table 1).

Table 1: Full model observations and (in parentheses) corrected values of first integrals

	$\langle H \rangle$	$\langle \lvert\boldsymbol{J}\rvert^2 \rangle$	$\langle H^2 \rangle$	$\langle \lvert\boldsymbol{J}\rvert^4 \rangle$	$\langle H \lvert\boldsymbol{J}\rvert^2 \rangle$
$H_{\text{full}} = -2$	-0.33 (-0.38)	4.59 (4.45)	0.22 (0.23)	-0.63 (-0.98)	34.58 (31.45)
$H_{\text{full}} = -1$	-0.11 (-0.18)	4.78 (4.68)	0.10 (0.12)	0.38 (-0.01)	37.55 (35.88)
$H_{\text{full}} = 0$	0.02 (-0.04)	4.63 (4.56)	0.08 (0.08)	0.90 (0.60)	35.26 (34.30)
$H_{\text{full}} = 1$	0.17 (0.15)	4.74 (4.75)	0.13 (0.12)	1.73 (1.61)	37.76 (37.44)
$H_{\text{full}} = 2$	0.31 (0.28)	4.87 (5.00)	0.22 (0.21)	2.49 (2.46)	39.26 (41.74)

Given the time averages, we compute the Lagrange multipliers using the algorithm described in Section 2.1 with prior distribution π the uniform distribution on the sphere. The Lagrange multipliers are also recorded in Table 2. The magnitude of $\gamma_{\{H,J,HJ\}}$ indicates that all observations are relevant for all but the most negative energy levels.

Subsequently, we run simulations of the corrected model using the computed parameters. Table 1 also records expectations from the thermostat-corrected model.

Table 2: Lagrange multipliers for each energy level.

	β_H	β_J	γ_H	γ_J	γ_{HJ}
$H_{\text{full}} = -2$	5.98	-0.20	0.69	0.41×10^{-3}	-0.04
$H_{\text{full}} = -1$	2.89	-0.03	2.67	9.77×10^{-3}	-0.33
$H_{\text{full}} = 0$	-0.76	0.20	3.38	9.97×10^{-3}	-0.37
$H_{\text{full}} = 1$	-3.54	0.37	4.29	15.31×10^{-3}	-0.54
$H_{\text{full}} = 2$	-6.42	0.53	4.45	14.05×10^{-3}	-0.51

By analogy with canonical statistical mechanics, we may think of the weak vortices that are ignored in the reduced model as forming a reservoir with which our reduced model exchanges energy and angular momentum. Experience with canonical statistical mechanics of point vortices in the plane [47, 14] suggests that for small reservoir sizes the canonical Gibbs distribution must be modified with higher order terms to agree with the full system statistics. Table 3 contains a study of the Lagrange multipliers as a function of the full system size M_{full}, confirming that the Lagrange multipliers γ_H, γ_J and γ_{HJ} are more significant for smaller M_{full}.

The energy of the strong vortices may become arbitrarily large because of the singularity in the

Table 3: Lagrange multipliers as a function of M_{full}, all for $H_{\text{full}} = 0$.

	β_H	β_J	γ_H	γ_J	γ_{HJ}
$M_{\mathbf{full}} = 36$	-1.51	1.35	27.75	117.15×10^{-3}	-3.07
$M_{\mathbf{full}} = 72$	-4.27	0.82	8.71	37.79×10^{-3}	-1.12
$M_{\mathbf{full}} = 144$	-0.97	0.32	6.70	25.80×10^{-3}	-0.82
$M_{\mathbf{full}} = 288$	-0.76	0.20	3.38	9.97×10^{-3}	-0.37
$M_{\mathbf{full}} = 576$	-1.09	0.13	0.87	3.08×10^{-3}	-0.10

Hamiltonian as two vortices approach each other. But the same is true for the energy in the reservoir. If there are at least three reservoir vortices and not all those vortices have the same sign, the reservoir can supply or remove any amount of energy.

The condition on the angular momentum is more interesting. The system of strong vortices, all with strength $\pm\Gamma_{\text{strong}}$, has angular momentum satisfying $|\boldsymbol{J}_{\text{red.}}| \leq M\Gamma_{\text{strong}}$. For the reservoir it holds that $|\boldsymbol{J}_{\text{weak}}| \leq (N - M)\Gamma_{\text{weak}}$. It is necessary that the reservoir can supply sufficient angular momentum, that is

$$M\Gamma_{\text{strong}} \leq (M_{\text{full}} - M)\Gamma_{\text{weak}} \Leftrightarrow \frac{M_{\text{full}} - M}{M} \geq \frac{\Gamma_{\text{strong}}}{\Gamma_{\text{weak}}}.$$

In the thermal bath simulations discussed in this section $M = 8$ and $\Gamma_{\text{strong}}/\Gamma_{\text{weak}} = 5$, this means M_{full} should satisfy $M_{\text{full}} \geq 48$. The smallest system considered ($M_{\text{full}} = 36$) does not, explaining its eccentric parameter values in Table 3.

4.1. Equilibrium results

In this section we compare statistical properties of the corrected model with those of the full and reduced models. In Figures 2–4 we show histograms of a number of solution features for the 8 vortex model: the distributions of H and $|\boldsymbol{J}|$, as well as typical distances between like- and opposite-signed vortices, a metric used by Bühler [47]. In each histogram, the statistics corresponding to the strong vortices in the full model, the reduced model, thermostat-corrected reduced model, and Metropolis-Hastings samples are displayed. Figures 2–4 correspond to approximate total energies $H_{\text{full}} \approx -2$, 0 and 2, respectively.

The full and reduced model simulations are performed with a time step of 5×10^{-3} and run up to $T = 5 \times 10^4$, taking 10^5 samples spaced evenly in time. For the Metropolis-Hastings method we use 10^6 samples. The same figures also show results from the thermostated system (dash-dot lines), run with a time step of 10^{-3} up to $T = 10^6$, taking 10^6 samples. The parameters in (7) were set to be $\varepsilon = 10$ and $\gamma = 0.1$. These results confirm that the thermostated system samples the least-biased density closely.

The reduced model is Hamiltonian and the Poisson integrator ensures that the energy is conserved with a standard deviation of order 10^{-3} and the angular momentum constant to machine precision. Both cases correspond to approximate delta-distributions in the upper histograms in Figures 2–4. Note that due to the high skewness of the distribution for $|\boldsymbol{J}|$, the observed mean differs significantly from the median and mode, implying some ambiguity in choosing the angular momentum for an appropriate initial condition for the reduced model.

A simple Hamiltonian reduced model is naturally incapable of sampling the energy and angular momentum spectra, since these quantities are first integrals. In turn, the reduced model shows

significant bias in statistics such as vortex separation. The thermostat-corrected model faithfully samples the least-biased probability density, as indicated by the good agreement in the histograms of the corrected model and Metropolis-Hastings samples. The least-biased density does a good job of approximating the strong-vortex statistics in the negative to moderate total energy regime. At large positive total energies, the strong vortex energy and angular momentum distributions are still well-represented by the least-biased PDF, but some bias in the vortex separations can be observed. The closeness of the thermostat results to those from the Metropolis-Hastings sampling indicate the error lies in the choice of least-biased density, not in the thermostat sampling.

4.2. Dynamic consistency

The results in the previous section confirm that the thermostated simulations lead to equilibrium distributions of observables H and $|\boldsymbol{J}|$ similar to those of the full system. In this section we address the issue of the degree to which our equilibrium correction mechanism disturbs dynamics, as encoded in autocorrelation functions and diffusivity. Diffusivity was considered by [48] for a system of identical point vortices and by [49] for a wide array of problems with scale separation. We emphasize that the values of the thermostat parameters ε and γ have no impact on the equilibrium statistics presented in the previous section, and only affect the rate at which the least-biased PDF is sampled. Faster convergence to the equilibrium distribution correlates with a larger deviation from the unperturbed dynamics and vice-versa.

4.2.1. Autocorrelation functions

Given a sequence of L equally spaced observation times $t_i \in [0, T]$ for $i \in [0, L]$, and the values of the relevant observable (in our case vortex position) $u_i = u(t_i)$ at those times, the discrete

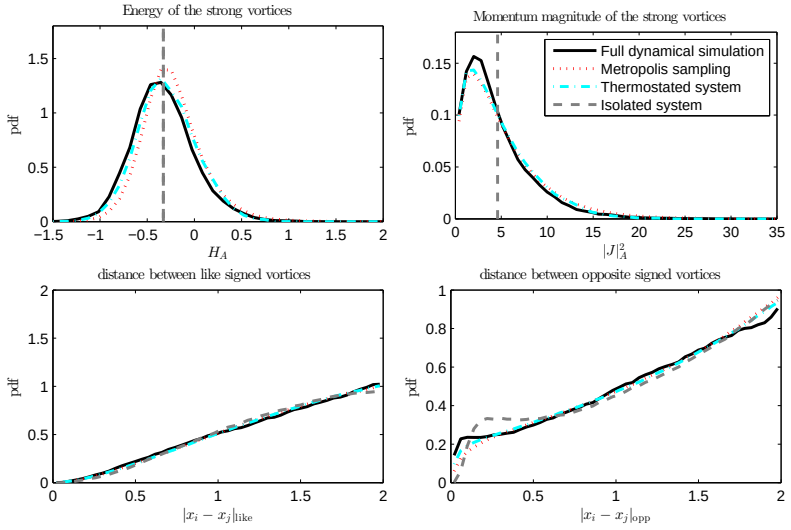

Figure 2: Histograms for $H_{\text{full}} \approx -2$. The upper left and right panels compare strong vortex energy and angular momentum magnitude. The lower left (resp. right) panel compares the distance between like (resp. opposite) signed strong vortices. The parameters are specified in the text.

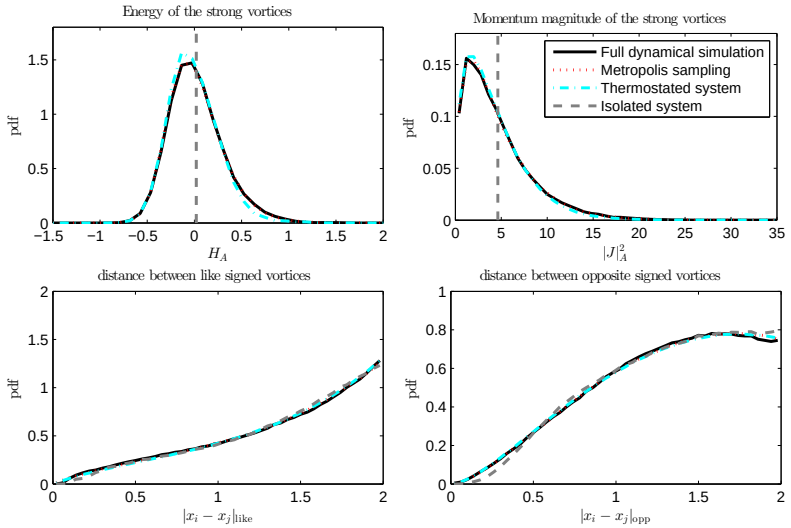

Figure 3: Histograms for $H_{\text{full}} \approx 0$. Same panel layout as Figure 2.

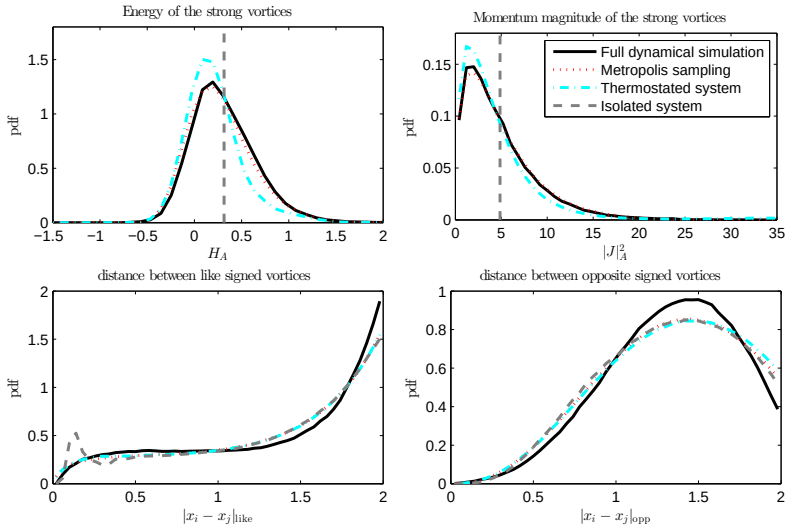

Figure 4: Histograms for $H_{\text{full}} \approx 2$. Same panel layout as Figure 2.

autocorrelation function is defined by

$$\nu_i^u = \frac{1}{L-i}\sum_{j=i}^{L} u(t_j)u(t_{j-i}).$$

A normalized autocorrelation function $\hat{\nu}^u$ is given by dividing each ν_i^u by ν_0^u, i.e. $\hat{\nu}_i^u = \nu_i^u/\nu_0^u$.

We average the autocorrelation function over all $3M$ (strong) vortex coordinates. Three symmetries in the problem justify this averaging: the vortex numbering is arbitrary; the choice of reference frame is arbitrary and the sign of the vortices appears in the dynamics as a reversal of time, to which the autocorrelation is insensitive. Additionally, the observables H and $|\boldsymbol{J}|$ are isotropic.

Furthermore we ensure that the phase space is well sampled by averaging the autocorrelation functions over an ensemble of P solutions. The choice of ensemble initial condition is detailed in Section 4.2.3 below. We then find the average autocorrelation function

$$\nu_i = \frac{1}{3MP(L-i)}\sum_{p=1}^{P}\sum_{m=1}^{M}\sum_{j=i}^{L} x_m^p(t_j)x_m^p(t_{j-i}) + y_m^p(t_j)y_m^p(t_{j-i}) + z_m^p(t_j)z_m^p(t_{j-i}) \tag{18}$$

and the normalized average autocorrelation function

$$\hat{\nu}_i = \frac{1}{MP(L-i)}\sum_{p=1}^{P}\sum_{m=1}^{M}\sum_{j=i}^{L} \boldsymbol{x}_m^p(t_j)\cdot\boldsymbol{x}_m^p(t_{j-i}), \tag{19}$$

where a superscript p represents the solution from ensemble member p. The normalized autocorrelation function in (19) follows from the Casimirs $C_i = \boldsymbol{x}_i(t)\cdot\boldsymbol{x}_i(t) = 1\,\forall\, i,t$.

4.2.2. Diffusivity

For general multiscale dynamical systems with a separation of slow and fast dynamics, it is often desirable to model fast forces by a diffusion process, resulting in stochastic differential equation of the form [50, 51]

$$dX = f(X)dt + K(X)dW,$$

where f represents the slow dynamics, W is a Wiener process and K is the diffusivity. The value of the diffusivity K can be estimated by sampling solutions to the original, multiscale, problem and applying Kubo's formula

$$K_{\Delta\tau} = \frac{\langle \Delta X \Delta X\rangle}{2\Delta\tau},$$

where ΔX represents displacement during the sampling interval $\Delta\tau$. Choosing the correct sampling interval is a notorious problem; for a comparison see [50].

If we take the average diffusivity for each vortex coordinate we find

$$\begin{aligned} K_{\Delta\tau} &= \frac{1}{6M\Delta\tau}\sum_{m=1}^{M}\langle \Delta x_m\Delta x_m + \Delta y_m\Delta y_m + \Delta z_m\Delta z_m\rangle \\ &= \frac{1}{6M\Delta\tau}\sum_{m=1}^{M}\langle \Delta\boldsymbol{x}_m\cdot\Delta\boldsymbol{x}_m\rangle\,. \end{aligned}$$

We assume the observations are given at the same times t_i as before and that the sampling time is an integer multiple of the observation interval, i.e. $\Delta\tau = i\frac{T}{L}$. With an ensemble of P simulations the diffusivity estimator would then be

$$\begin{aligned} K_{\Delta\tau} &= \frac{1}{6MP\Delta\tau}\sum_{p=1}^{P}\sum_{m=1}^{M}(\boldsymbol{x}_m^p(t_i)-\boldsymbol{x}_m^p(0))\cdot(\boldsymbol{x}_m^p(t_i)-\boldsymbol{x}_m^p(0)) \\ &= \frac{1}{6MP\Delta\tau}\sum_{p=1}^{P}\sum_{m=1}^{M}2-2\boldsymbol{x}_m^p(0)\cdot\boldsymbol{x}_m^p(t_i) \\ &= \frac{1}{3\Delta\tau}-\frac{1}{3MP\Delta\tau}\sum_{p=1}^{P}\sum_{m=1}^{M}\boldsymbol{x}_m^p(0)\cdot\boldsymbol{x}_m^p(t_i), \end{aligned}$$

where again a superscript p denotes the solution from ensemble member p. Averaging over all time series data yields the estimator:

$$\begin{aligned} K_{\Delta\tau}^i &= \frac{1}{3\Delta\tau}-\frac{1}{3MP(L-i)\Delta\tau}\sum_{p=1}^{P}\sum_{j=i}^{L}\sum_{m=1}^{M}\boldsymbol{x}_m^p(t_{j-i})\cdot\boldsymbol{x}_m^p(t_j) \\ &= \frac{1-\hat{\nu}_i}{3\Delta\tau}. \end{aligned}$$

This *shift-averaged estimator* is shown in [49] to improve the quality of the estimator.

4.2.3. Dynamical results

In Figure 5 we compare auto-correlation functions for the strong vortices in the full and reduced models as well as for the thermostat-corrected model over a range of parameters ε and γ. The thick solid black line represents the result for an (unthermostated) system in contact with 280 weak ($\Gamma_B = \pm\frac{1}{5}$) vortices, with a total energy $H_{\text{full}} = 0$. The results present the average over an ensemble of 1000 runs. For each simulation the initial placement of each strong vortex was taken uniformly over the sphere and the weak vortices were placed such that the full system satisfied $H_{\text{full}} = 0$ and $\boldsymbol{J}_{\text{full}} = 0$. The thick dashed black line represents the results for an ensemble of simulations of the isolated system, with everything else unchanged. The other lines represent results for thermostated simulations using the parameters as given in Table 2 for the case of $H = 0$.

The corresponding diffusivity constants are presented in Figure 6. The results are taken from the same simulations as described in the paragraph above. Because this figure is visually more striking, we shall limit our discussion to the diffusion constant.

For ε small, the thermostat perturbation is weak, and both autocorrelation functions and diffusivity approach those of the reduced model with constant H, $\boldsymbol{J}$. Also, the autocorrelations are insensitive to the parameter γ in this regime. For even smaller ε the autocorrelations and diffusivities become indistinguishable from those of the reduced model. Hence even though the dynamics samples the least-biased density on long time scales, its short time dynamics is similar to an unperturbed model. For moderate ε, dependence on γ becomes more pronounced, and a diffusivity closer to that of the full model can be achieved. For even larger values of ε, the diffusivity becomes much more sensitive to the value of γ, as indicated in Figure 6c.

Figure 6d has been included to illustrate two important properties. Firstly, as the sampling interval goes to zero, the estimator of the diffusivity constant shows linear behavior. This is in

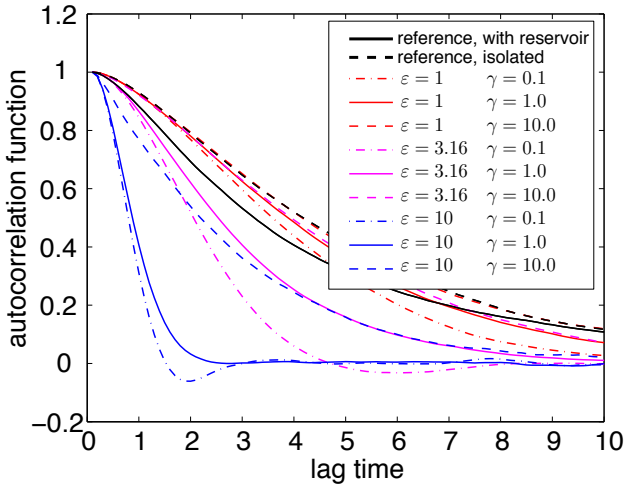
autocorrelation function
lag time
reference, with reservoir
reference, isolated
ε = 1 γ = 0.1
ε = 1 γ = 1.0
ε = 1 γ = 10.0
ε = 3.16 γ = 0.1
ε = 3.16 γ = 1.0
ε = 3.16 γ = 10.0
ε = 10 γ = 0.1
ε = 10 γ = 1.0
ε = 10 γ = 10.0

Figure 5: Comparison of autocorrelation of the vortex coordinates. The bold lines are two reference cases: the full model (solid) and the reduced model (dashed). The thin lines indicate autocorrelation functions of the thermostat-corrected model for indicated values of parameters ε and γ.

agreement with known results for the GBK thermostat[52] and is an improvement on Langevin thermostats, which would erroneously tend to a constant value as the sampling interval is decreased. Secondly, for large sampling interval the estimator shows an inverse linear tendency. This corresponds simply to the decorrelation of the vortex dynamics.

4.3. Adaptive determination of multipliers

Consider the same reduced point vortex model of 8 vortices with $\Gamma = \pm 1$ and assume observations on the energy and momentum are known from a simulation of the full system including the thermal bath. We start such a simulation with an ensemble of $P = 100$ initial conditions drawn from the uniform prior. The time step is chosen as 1×10^{-2} and the method described in Section 2.1 for updating the Lagrange multipliers is applied every time unit, i.e. $M = 100$. Between subsequent updates of the multipliers, the maximum difference is limited by $|\Delta\lambda_k| \leq 0.1$. When using equilibrium statistics, this limit only affects the beginning of the simulation, when the small sample size used leads to a large variance in the estimators.

The target observation values are taken from a simulation of strong vortices interacting with a thermal bath of weak vortices. Three different averages are used.

1. In Figure 7 the long time mean is taken and used throughout.
2. In Figure 8 the running mean is used. This reflects the situation where we have no a priori knowledge of the observations, and are continuously feeding new real-time data into the simulation.
3. In Figure 9 a time-localized average of the observable is used. The averaging has a time-scale of a 100 time units. This also corresponds to feeding the simulation new data, but now the assumption of equilibrium is relaxed.

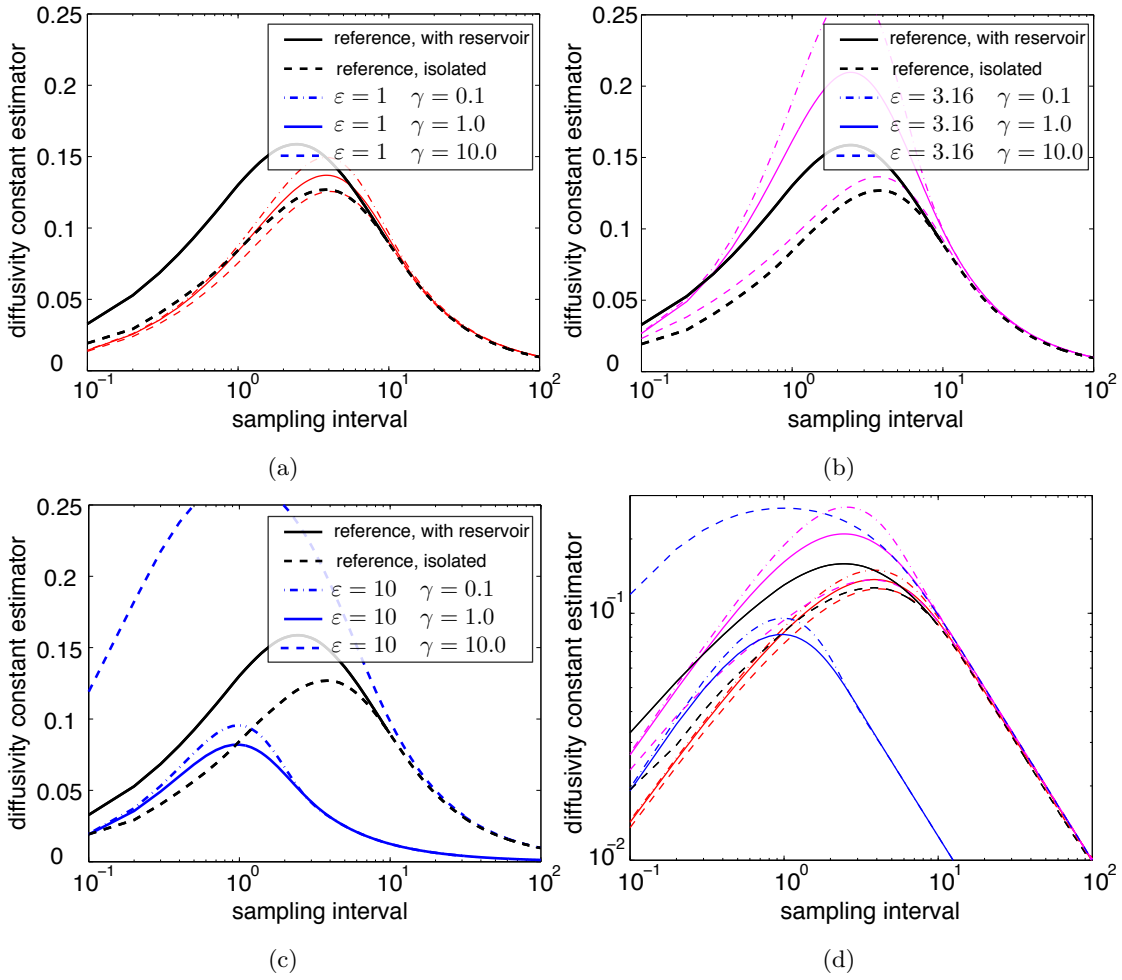
diffusivity constant estimator
sampling interval
reference, with reservoir
reference, isolated
$\varepsilon = 1$ $\gamma = 0.1$
$\varepsilon = 1$ $\gamma = 1.0$
$\varepsilon = 1$ $\gamma = 10.0$
(a)
$\varepsilon = 3.16$ $\gamma = 0.1$
$\varepsilon = 3.16$ $\gamma = 1.0$
$\varepsilon = 3.16$ $\gamma = 10.0$
(b)
$\varepsilon = 10$ $\gamma = 0.1$
$\varepsilon = 10$ $\gamma = 1.0$
$\varepsilon = 10$ $\gamma = 10.0$
(c)
(d)

Figure 6: Comparison of average diffusivity constant as a function of sampling intervals. In all figures, the bold lines indicate two reference cases: the full model (solid) and the reduced model (dashed). Subfigures (a), (b) and (c) show thermostated simulation results for ε equal to 10^0, $10^{0.5}$ and 10^1 respectively. The value of γ is represented by dash-dot (10^{-1}), solid (10^0) or dashed (10^1) lines. A combined log-log plot of all parameter values is given in subfigure (d).

When using either a long time mean observation or a running mean observation, the simulation results tends towards the correct long-time averages. When using time-local averages the simulation averages appear to tend towards a similar value. In all three cases the instantaneous ensemble mean remains close to the (moving) target for both energy and momentum. This is especially notable for the third case, where the target varies over time, but the simulation ensemble mean follows closely, with only a little lag.

The inaccuracies incurred during the first approximately 100 time units indicate that the prior does not match the observed state well. This results in the (negative) growth of β_H being limited briefly at the beginning of each simulation. Subsequently, both Lagrange multipliers appear to oscillate irregularly about some mean value for the first two cases. In the case of a shifting target,

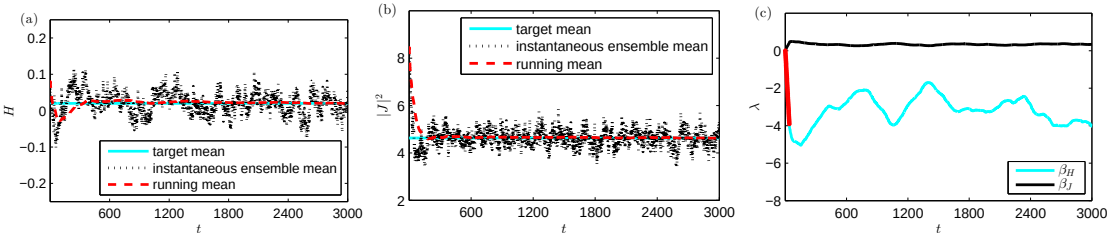

Figure 7: Results when using *long-time mean* observations as a target while adaptively determing the Lagrange multipliers. Target observations for Hamiltonian (a) and momentum magnitude (b) are overlaid with the instantaneous ensemble mean (black dotted) and the running ensemble mean (red solid) from simulation. (c): Lagrange multipliers, the red dots indicate time steps at which their rate of change was limited.

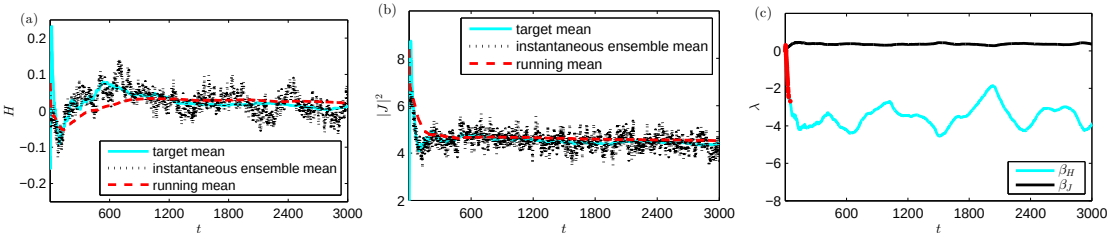

Figure 8: Results when using *running mean* observations as a target while adaptively determing the Lagrange multipliers. Target observations for Hamiltonian (a) and momentum magnitude (b) are overlaid with the instantaneous ensemble mean (black dotted) and the running ensemble mean (red solid) from simulation. (c): Lagrange multipliers, the red dots indicate time steps at which their rate of change was limited.

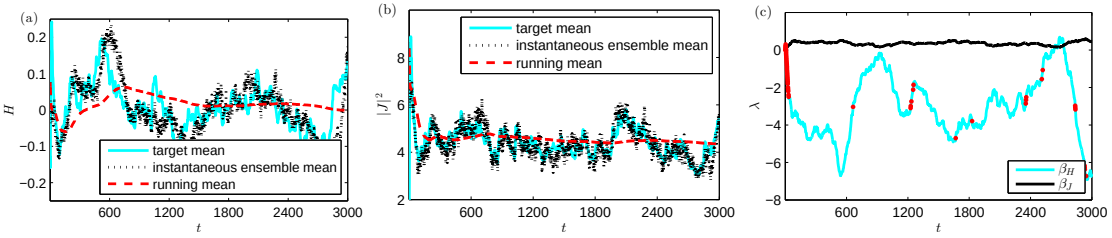

Figure 9: Results when using *time-local averaged* observations as a target while adaptively determing the Lagrange multipliers. Target observations for Hamiltonian (a) and momentum magnitude (b) are overlaid with the instantaneous ensemble mean (black dotted) and the running ensemble mean (red solid) from simulation. (c): Lagrange multipliers, the red dots indicate time steps at which their rate of change was limited.

the Lagrange multipliers vary in time more erratically, as is to be expected. This results in the limiter being active for a few brief periods of the simulation.

5. Conclusion

In this article we propose a method for perturbing trajectories of numerical simulations to correct for equilibrium observations. We treat the restricted case of a Hamiltonian ODE with observations on the set of first integrals of the system. The approach entails perturbing the solutions using a stochastic thermostat such that they become ergodic in a prescribed invariant measure: the least-biased density corresponding to a maximum entropy treatment of the observations.

We apply the approach to the case of model reduction in a heterogeneous system of weak and strong point vortices on a sphere, in which observations of the energy H and angular momentum magnitude $|\boldsymbol{J}|$ are made on a subsystem consisting of the strong vortices. A reduced model is constructed by neglecting the weak vortices, and the expectations of the reduced model are corrected using the proposed methodology.

Numerical experiments confirm that the distributions of the observed quantities H and $|\boldsymbol{J}|$ can be well represented using the thermostat technique. Other equilibrium metrics such as the distribution of distances between like- and opposite-signed vortices are also in agreement across a range of total energy values of the full system, although some discrepancies occur at large positive energies.

We also investigated the degree to which correction of trajectories for expectations may affect dynamical information in the form of autocorrelation functions and diffusivity. By decreasing the perturbation parameter ε of the thermostat, the autocorrelation functions of the unperturbed, reduced system may be precisely recovered. As ε is increased, one may increase the diffusivity to values that agree with the full system. This is consistent with results reported in [52] in the context of molecular dynamics where it was shown that the GBK thermostat used here approaches Langevin dynamics in the limit of large stochastic forcing.

6. Acknowledgements

The work of the first author was supported by a grant financed by the Netherlands Organisation for Scientific Research (NWO), as part of research programme 613.001.009. The third author was supported by grant EP/G036136/1 from the Engineering and Physical Sciences Research Council (UK).

Appendix A. Integration of the two-vortex system

Each pair flow $\phi_{\Delta t}^{i,j}$ is the flow associated with the two-vortex problem $\dot{y} = B\nabla H_{ij}$ for vortex pair (i, j) given by

$$\dot{\boldsymbol{x}}_i = \frac{-\Gamma_j}{4\pi} \cdot \frac{\boldsymbol{x}_j \times \boldsymbol{x}_i}{1 - \boldsymbol{x}_i \cdot \boldsymbol{x}_j}, \tag{A.1}$$

$$\dot{\boldsymbol{x}}_j = \frac{-\Gamma_i}{4\pi} \cdot \frac{\boldsymbol{x}_i \times \boldsymbol{x}_j}{1 - \boldsymbol{x}_i \cdot \boldsymbol{x}_j}. \tag{A.2}$$

This is again a Poisson system, the particular case of (10) with $M = 2$. It has a Hamiltonian $H_{ij} = \frac{\Gamma_i \Gamma_j}{8\pi} \ln(2 - 2\boldsymbol{x}_i \cdot \boldsymbol{x}_j)$ and angular momentum $\boldsymbol{J}_{ij} = \Gamma_i \boldsymbol{x}_i + \Gamma_j \boldsymbol{x}_j$.

Observe that due to the conservation of H_{ij} the denominators in (A.1)–(A.2) are constant. In the numerators of (A.1)–(A.2) we may subsitute $\Gamma_j \boldsymbol{x}_j = \boldsymbol{J}_{ij} - \Gamma_i \boldsymbol{x}_i$ and use $\boldsymbol{x}_i \times \boldsymbol{x}_i = 0$ to find

$$\dot{\boldsymbol{x}}_i = \frac{-1}{4\pi} \cdot \frac{\boldsymbol{J}_{ij}}{1 - \boldsymbol{x}_i \cdot \boldsymbol{x}_j} \times \boldsymbol{x}_i = \widehat{\boldsymbol{a}} \boldsymbol{x}_i, \tag{A.3}$$

$$\dot{\boldsymbol{x}}_j = \frac{-1}{4\pi} \cdot \frac{\boldsymbol{J}_{ij}}{1 - \boldsymbol{x}_i \cdot \boldsymbol{x}_j} \times \boldsymbol{x}_j = \widehat{\boldsymbol{a}} \boldsymbol{x}_j, \tag{A.4}$$

where $\boldsymbol{a} = \frac{-1}{4\pi} \frac{\boldsymbol{J}}{1 - \boldsymbol{x}_i \cdot \boldsymbol{x}_j}$ and $\widehat{\boldsymbol{a}}$ denotes the skew-symmetric matrix associated to the cross product with a vector $\boldsymbol{a} = (a_1, a_2, a_3)^T$ by

$$\widehat{\boldsymbol{a}} = \begin{bmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{bmatrix}.$$

The structure of $\widehat{\boldsymbol{a}}$ admits efficient evaluation of the solution $\boldsymbol{x}(t) = \exp(\widehat{\boldsymbol{a}}t)\boldsymbol{x}(0)$ using Rodigues's formula (see, e.g. [7])

$$\exp(\widehat{\boldsymbol{a}}t) = I + \frac{\sin(\alpha t)}{\alpha} \widehat{\boldsymbol{a}} + \frac{1 - \cos(\alpha t)}{\alpha^2} \widehat{\boldsymbol{a}}^2,$$

where $\alpha = |\boldsymbol{a}|$. This is a computationally efficient expression for the given form of $\widehat{\boldsymbol{a}}$ because the full matrix exponential need not be computed. Instead we introduce $\tilde{\boldsymbol{a}} = \boldsymbol{a}/\alpha$ and write

$$\exp(\widehat{\boldsymbol{a}}t)\boldsymbol{x}_0 = \boldsymbol{x}_0 + \sin(\alpha t)\tilde{\boldsymbol{a}} \times \boldsymbol{x}_0 + (1 - \cos(\alpha t))\left(\tilde{\boldsymbol{a}}\left(\tilde{\boldsymbol{a}} \cdot \boldsymbol{x}_0\right) - \boldsymbol{x}_0\right).$$

Appendix B. Time integration of the thermostated system

In this section we detail the method of integration for the perturbation dynamics and thermostat variable.

Just as the unperturbed dynamics, the modified double bracket thermostat is a combination of pairwise vortex interactions. Hence it is natural to apply the same splitting to g as was applied to f. In fact, the pairwise thermostat interaction can be executed along with the unperturbed dynamics. Recall from Section 3.2 that we denote by $\phi^{i,j}_{\Delta t}$ the time Δt flow map associated with f_{ij}. Similarly we define by $\gamma^{i,j}_{\Delta t}$ the time Δt flow map associated with the perturbation dynamics of a vortex pair (i, j). The flow map of the dynamics of the thermostat variable ξ is represented by $\chi_{\Delta t}$. A symmetric composition of these flows is given by

$$\Phi_{\Delta t} = \prod_{(i,j) \in C} \left(\phi^{i,j}_{\Delta t/2} \circ \gamma^{i,j}_{\Delta t/2}\right) \circ \chi_{\Delta t} \circ \prod_{(i,j) \in C^*} \left(\gamma^{i,j}_{\Delta t/2} \circ \phi^{i,j}_{\Delta t/2}\right).$$

While advancing ξ by $\chi_{\Delta t}$, it is assumed that y is fixed and therefore $h(y)$ is a constant. This means the dynamics of $\xi(t)$ is just an Ornstein-Uhlenbeck process with mean $\frac{h}{\gamma}$ and unit variance. The exact solution is given by

$$\chi_{\Delta t}\, \xi_0 = \xi_0 e^{-\gamma \Delta t} + \frac{h}{\gamma}(1 - e^{-\gamma \Delta t}) + e^{-\gamma \Delta t} W(e^{2\gamma \Delta t} - 1).$$

The integration of the perturbed dynamics for a vortex pair (i, j) means integrating the system of ordinary differential equations given by

$$\dot{\boldsymbol{x}}_i = \frac{\Gamma_j}{4\pi}\boldsymbol{x}_i \times \boldsymbol{x}_i \times \boldsymbol{x}_j = \frac{\Gamma_i + \Gamma_j}{4\pi}\boldsymbol{x}_i \times \boldsymbol{x}_i \times \boldsymbol{x}_j + \frac{\Gamma_i - \Gamma_j}{4\pi}\boldsymbol{x}_i \times \boldsymbol{x}_i \times \boldsymbol{x}_j,$$
$$\dot{\boldsymbol{x}}_j = \frac{\Gamma_i}{4\pi}\boldsymbol{x}_j \times \boldsymbol{x}_j \times \boldsymbol{x}_i = \frac{\Gamma_i + \Gamma_j}{4\pi}\boldsymbol{x}_j \times \boldsymbol{x}_j \times \boldsymbol{x}_i - \frac{\Gamma_i - \Gamma_j}{4\pi}\boldsymbol{x}_j \times \boldsymbol{x}_j \times \boldsymbol{x}_i.$$

The equations have been split into a symmetric and an anti-symmetric part. A symmetric composition may once more be applied to integrate the two parts. In our case however, the thermostated system will consist only of vortices with strength $\pm\Gamma_{\text{strong}}$ and thus each vortex pair interaction is either fully symmetric or fully antisymmetric.

By introducing $\alpha_S = \frac{\Gamma_i+\Gamma_j}{4\pi}$ the symmetric part may be written as

$$\dot{\boldsymbol{x}}_i = \alpha_S \boldsymbol{x}_i \times \boldsymbol{x}_i \times \boldsymbol{x}_j = \alpha_S \left(\boldsymbol{x}_j \times \boldsymbol{x}_i\right) \times \boldsymbol{x}_i, \tag{B.1}$$
$$\dot{\boldsymbol{x}}_j = \alpha_S \boldsymbol{x}_j \times \boldsymbol{x}_j \times \boldsymbol{x}_i = -\alpha_S \left(\boldsymbol{x}_j \times \boldsymbol{x}_i\right) \times \boldsymbol{x}_j. \tag{B.2}$$

These dynamics are symmetric with respect to the plane equidistant between the two vortices. This is because the dynamics are rotations in opposite direction about the same vector $\boldsymbol{x}_j \times \boldsymbol{x}_i$ and this vector must lie in said plane. Furthermore the dynamics of (B.1)–(B.2) does not allow the vortex pair to pass through the position where they are antipedian, as in this case their cross product is zero. All in all this means the final position of the two vortices is uniquely determined by their chord distance.

The change in the distance between the two vortices can be represented by the change of their inner product

$$\begin{aligned}\frac{\partial}{\partial t}\left(\boldsymbol{x}_i \cdot \boldsymbol{x}_j\right) &= \dot{\boldsymbol{x}}_i \cdot \boldsymbol{x}_j + \boldsymbol{x}_i \cdot \dot{\boldsymbol{x}}_j \\ &= 2\alpha_S \left(\left(\boldsymbol{x}_i \cdot \boldsymbol{x}_j\right)^2 - 1\right)\end{aligned}$$

The solution to this differential equation is given by

$$\boldsymbol{x}_i \cdot \boldsymbol{x}_j|_{t=\Delta t} = \tanh\left(2\alpha\Delta t - \operatorname{artanh}\left(\boldsymbol{x}_i \cdot \boldsymbol{x}_j\right)|_{t=0}\right).$$

In the anti-symmetric part we introduce $\alpha_A = \frac{\Gamma_i-\Gamma_j}{4\pi}$ to write

$$\dot{\boldsymbol{x}}_i = \alpha_A \boldsymbol{x}_i \times \boldsymbol{x}_i \times \boldsymbol{x}_j,$$
$$\dot{\boldsymbol{x}}_j = -\alpha_A \boldsymbol{x}_j \times \boldsymbol{x}_j \times \boldsymbol{x}_i.$$

By rearranging the order of the cross product we achieve

$$\dot{\boldsymbol{x}}_i = \alpha_A \left(\boldsymbol{x}_j \times \boldsymbol{x}_i\right) \times \boldsymbol{x}_i = \widehat{a}_A \boldsymbol{x}_i, \tag{B.3}$$
$$\dot{\boldsymbol{x}}_j = \alpha_A \left(\boldsymbol{x}_j \times \boldsymbol{x}_i\right) \times \boldsymbol{x}_j = \widehat{a}_A \boldsymbol{x}_j. \tag{B.4}$$

The vector $\boldsymbol{a}_A$ implicity defined by (B.3)–(B.4) can simply be shown to be constant under the antisymmetric flow. This means the dynamics of the anti-symmetric part may be integrated by using Rodigues's formula, similar to the unperturbed dynamics (A.3)–(A.4).

References

[1] M. Allen, D. Tildesley, Computer Simulation of Liquids, Oxford University Press, 1989, iSBN:978-0198556459.

[2] T. Schlick, Molecular Modeling and Simulation: An Interdisciplinary Guide (Interdisciplinary Applied Mathematics), 2nd Edition, Springer, 2010, iSBN:978-1441963505.

[3] P. Holmes, J. L. Lumley, G. Berkooz, Turbulence, coherent structures, dynamical systems and symmetry, Cambridge University Press, 1998.

[4] P. A. Davidson, Turbulence: An Introduction for Scientists and Engineers, Oxford University Press, 2004.

[5] W. Hayes, K. R. Jackson, A fast shadowing algorithm for high dimensional ODE systems, SIAM Journal on Scientific Computing 29 (4) (2007) 1738–1758.

[6] E. Hairer, C. Lubich, G. Wanner, Geometric Numerical Integration, 2nd Edition, Springer Series in Computational Mathematics, Springer-Verlag, New York, 2006.

[7] B. Leimkuhler, S. Reich, Simulating Hamiltonian Dynamics, 1st Edition, Camb. Monogr. on Appl. and Comput. Math., Camb. Univ. Press, Camb., U.K., 2004.

[8] R. V. Abramov, A. J. Majda, Statistically relevant conserved quantities for truncated quasigeostrophic flow, Proceedings of the National Academy of Sciences 100 (7) (2003) 3841–3846.

[9] S. Dubinkina, J. Frank, Statistical mechanics of Arakawa's discretizations, Journal of Computational Physics 227 (2) (2007) 1286 – 1305. doi:http://dx.doi.org/10.1016/j.jcp.2007.09.002.
URL `http://www.sciencedirect.com/science/article/pii/S0021999107003944`

[10] S. Dubinkina, J. Frank, Statistical relevance of vorticity conservation in the hamiltonian particle-mesh method, Journal of Computational Physics 229 (7) (2010) 2634–2648.

[11] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, E. Teller, Equation of state calculations by fast computing machines, The Journal of Chemical Physics 21 (6) (1953) 1087–1092.

[12] E. J. Barth, B. B. Laird, B. J. Leimkuhler, Generating generalized distributions from dynamical simulation, The Journal of Chemical Physics 118 (13) (2003) 5759–5768.

[13] J. Bajars, J. Frank, B. Leimkuhler, Weakly coupled heat bath models for Gibbs-like invariant states in nonlinear wave equations, Nonlinearity 26 (7) (2013) 1945.

[14] S. Dubinkina, J. Frank, B. Leimkuhler, Simplified modelling of a thermal bath system, with application to a fluid vortex system, Multiscale Model Simul. 8 (5) (2010) 1882–1901.

[15] N. Ding, Y. Fang, R. Babbush, C. Chen, R. D. Skeel, H. Neven, Bayesian sampling using stochastic gradient thermostats, in: Advances in Neural Information Processing Systems (NIPS), 2014.

[16] E. T. Jaynes, Information theory and statistical mechanics, Physical review 106 (4) (1957) 620.

[17] E. T. Jaynes, Information theory and statistical mechanics. II, Phys. Rev. 108 (1957) 171–190. doi:10.1103/PhysRev.108.171.
URL http://link.aps.org/doi/10.1103/PhysRev.108.171

[18] N. Agmon, Y. Alhassid, R. D. Levine, An algorithm for finding the distribution of maximal entropy, Journal of Computational Physics 30 (2) (1979) 250–258.

[19] H. Haken, Information and Self-Organization: A Macroscopic Approach to Complex Systems, Springer, 2006.

[20] S. Davis, G. Gutiérrez, Conjugate variables in continuous maximum-entropy inference, Physical Review E 86 (5) (2012) 051136.

[21] A. J. Majda, B. Gershgorin, Improving model fidelity and sensitivity for complex systems through empirical information theory, Proceedings of the National Academy of Sciences 108 (25) (2011) 10044–10049.

[22] R. C. Dewar, C. H. Lineweaver, R. K. Niven, K. Regenauer-Lieb (Eds.), Beyond the Second Law: Entropy Production and Non-equilibrium Systems, Understanding Complex Systems, Springer, 2014.

[23] A. J. Majda, X. Wang, Nonlinear Dynamics and Statistical Theories for Basic Geophysical Flows, 1st Edition, Camb. Univ. Press, Camb., UK, 2006.

[24] J. C. Mattingly, A. Stuart, D. D.J. Higham, Ergodicity for sdes and approximations: locally lipschitz vector fields and degenerate noise, Stochastic processes and their applications 101 (2) (2002) 185–232.

[25] T. DelSole, A fundamental limitation of Markov models, Journal of the Atmospheric Sciences 57 (13) (2000) 2158–2168.

[26] S. Nosé, A molecular dynamics method for simulations in the canonical ensemble, Molecular physics 52 (2) (1984) 255–268.

[27] S. Nosé, A unified formulation of the constant temperature molecular dynamics methods, The Journal of Chemical Physics 81 (1) (1984) 511–519.

[28] W. G. Hoover, Canonical dynamics: equilibrium phase-space distributions, Physical Review A 31 (3) (1985) 1695.

[29] A. A. Samoletov, C. P. Dettmann, M. A. Chaplain, Thermostats for "slow" configurational modes, Journal of Statistical Physics 128 (6) (2007) 1321–1336.

[30] B. Leimkuhler, E. Noorizadeh, F. Theil, A gentle thermostat for molecular dynamics, J. Stat. Phys. 135 (2) (2009) 261–277.

[31] B. Leimkuhler, E. Noorizadeh, O. Penrose, Comparing the efficiencies of stochastic isothermal molecular dynamics models, J. Stat. Phys. 143 (5) (2011) 921–942.

[32] B. Leimkuhler, Generalized Bulgac-Kusnezov methods for sampling of the Gibbs-Boltzmann measure, Physical Review E 81 (2) (2010) 026703.

[33] A. Bulgac, D. Kusnezov, Canonical ensemble averages from pseudomicrocanonical dynamics, Phys. Rev. A 42 (1990) 5045–5048. doi:10.1103/PhysRevA.42.5045.
URL http://link.aps.org/doi/10.1103/PhysRevA.42.5045

[34] J. Bajars, J. Frank, B. Leimkuhler, Stochastic-dynamical thermostats for constraints and stiff restraints, Eur. Phys. J. Special Topics 200 (2011) 131–152.

[35] P. K. Newton, The N-Vortex Problem: Analytical Techniques, 1st Edition, Springer-Verlag, New York, 2001.

[36] A. J. Majda, A. L. Bertozzi, Vorticity and Incompressible Flow, 1st Edition, Camb. Texts in Appl. Math., Camb. Univ. Press, Camb., UK, 2002.

[37] G.-H. Cottet, P. D. Koumoutsakos, Vortex Methods, 1st Edition, Camb. Univ. Press, Camb., UK, 2000.

[38] L. Faddeev, L. Takhtajan, Hamiltonian Methods in the Theory of Solitons, Springer Series in Soviet Mathematics, Springer-Verlag, New York, 1987.

[39] J. Frank, Geometric space-time integration of ferromagnetic materials, Appl. Numer. Math. 48 (3-4) (2004) 307–322. doi:10.1016/j.apnum.2003.11.003.
URL http://dx.doi.org/10.1016/j.apnum.2003.11.003

[40] J. Frank, W. Huang, B. Leimkuhler, Geometric integrators for classical spin systems, J. Comput. Phys. 133 (1) (1997) 160–172. doi:10.1006/jcph.1997.5672.
URL http://dx.doi.org/10.1006/jcph.1997.5672

[41] S. Kehrbaum, J. H. Maddocks, Elasitc rods, rigid bodies, quaternions and the last quadrature, Phil. Trans. R. Soc. Lond. A 355 (1732) (1997) 2117–2136.

[42] O. Buhler, A Brief Introduction to Classical, Statistical and Quantum Mechanics, 1st Edition, Courant Lect. Notes, Am. Math. Soc./Courant Inst. of Math. Sci., Providence, R.I., 2006.

[43] M.-Q. Zhang, M.-Z. Qin, Explicit symplectic schemes to solve vortex systems, Computers Math. Applic. 26 (5) (1993) 51–56.

[44] G. W. Patrick, Dynamics of perturbed relative equilibria of point vortices on the sphere or plane, Nonlinear Science 10 (2000) 401–415.

[45] A. Bloch, P. Krishnaprasad, J. Marsden, T. Ratiu, The Euler-Poincaré equations and double bracket dissipation, Communications in Mathematical Physics 175 (1) (1996) 1–42.

[46] L. Rey-Bellet, Ergodic properties of Markov processes, Open Quantum Systems II (2006) 1–39.

[47] O. Bühler, Statistical mechanics of strong and weak point vortices in a cylinder., Phys. of Fluids 14 (7) (2002) 2139–2149.

[48] P. Chavanis, Kinetic theory of point vortices: Diffusion coefficient and systematic drift, Physical Review E 64 (2) (2001) 026309.

[49] C. Cotter, G. Pavliotis, Estimating eddy diffusivities from noisy Lagrangian observations, Communications in Mathematical Sciences 7 (4) (2009) 805–838.

[50] G. Pavliotis, A. Stuart, Multiscale Methods: Averaging and Homogenization, Springer, 2008.

[51] U. Hornung, Homogenization and Porous Media, Vol. 6, Springer Verlag, 1997.

[52] J. Frank, G. Gottwald, The Langevin limit of the Nosé-Hoover-Langevin thermostat, Journal of Statistical Physics 143 (4) (2011) 715–724.